

EUROPEAN PATENT APPLICATION

Application number: 89304816.5

Int. Cl.⁵: G06F 15/403, G06F 9/44

Date of filing: 12.05.89

The title of the invention has been amended
(Guidelines for Examination in the EPO, A-III,
7.3).

Date of publication of application:
14.11.90 Bulletin 90/46

Designated Contracting States:
DE FR GB

Applicant: International Business Machines
Corporation
Old Orchard Road
Armonk, N.Y. 10504(US)

Inventor: Macfarlane, Peter Robert
16 Fordington Avenue
Winchester Hampshire SO22 5AW(GB)

Representative: Bailey, Geoffrey Alan
IBM United Kingdom Limited Intellectual
Property Department Hursley Park
Winchester Hampshire SO21 2JN(GB)

Version-control in a data-processing system.

A data processing system having at least one local processor storing a version of at least one file and a remote processor determining from information indicative of said version stored by said local processor if said version stored by said local processor should be replaced by a version of said at least one file stored by said remote processor, characterised in that upon receipt of said replacement version said local processor determines if any intervening change has occurred in said version stored by said local processor since said information was supplied to said remote processor and allows replacement of said file if no intervening change has occurred. The system recognises and provides the solution to a number of problems that can arise in a system in which a local processor is automatically updated by more than one remote processor. This system recognises if another host processor has already updated the local processor and prevents a second possibly erroneous update taking place.

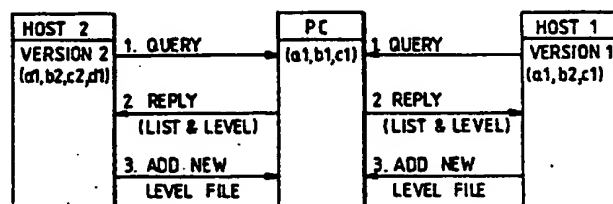


FIG. 3

DATA PROCESSING SYSTEM

This invention relates to the field of data processing systems. More particularly this invention relates to data processing systems with both local and remote processors.

It is known to provide data processing systems in which one or more workstations having a local processor and storing one or more files are connected to a remote processor, such as a mainframe computer. The term file is used to indicate any information such as application programs, operating systems, databases, dynamic data, etc, which is or may be held on a data processing system. A problem with such systems is that the files stored by a workstation often require updating so as to repair defects in software, update database material etc.

Europeans Published Patent Application No. 284924 describes a system in which a workstation supplies a mainframe with a list of what versions of what files are stored by the workstation upon a request from the workstation. The mainframe compares the versions of the files stored by the workstation with the versions of those files stored by the mainframe and automatically sends the workstation replacement updated versions of the files where necessary.

This system overcomes some of the problems associated with the maintenance of files on workstations by providing a system in which the updating of the files takes place automatically without the need for any user actions.

Viewed from one aspect the present invention provides a data processing system having at least one local processor storing a version of at least one file and a remote processor determining from information of said version stored by said local processor if said version stored by said local processor should be replaced by a version of said one file stored by said remote processor, characterised in that upon receipt of said replacement version said local processor determines if any intervening change has occurred in said version stored by said local processor since said information was supplied to said remote processor and allows replacement of said file if no intervening change has occurred.

The present invention both recognises and solves a problem which can occur when seeking to provide a system which includes the feature of automatically updating files stored on local processors by those stored by remote processors. The problem is illustrated by considering a system in which a local processor is connected to more than one remote processor. In this case the local processor may provide an indication of the files it has stored to more than one remote processor with the

result that a number of different remote processors may then compete to update the files stored by the local processor.

Such a system could result in a file which had already been updated by another remote processor being replaced again unnecessarily. The present invention provides a system in which the local processor checks to see that there has not been an intervening update (i.e. between the remote processor receiving the information as to what files are held on the local processor and the local processor receiving the replacement file) of the file in question before allowing the update to proceed.

Any intervening update could be the result of the action of another remote processor or possibly the result of action by the operator of the local processor. It will however be seen that the invention is particularly applicable to systems having a plurality of remote processors since it overcomes the problem of potential conflicts between the remote processors.

In preferred embodiments of the present invention said local processor supplies said information to said remote processor. This feature has the advantage that the remote processor need not maintain records of what versions of what files are stored by said remote processor. This lessens the storage and processing overhead that must be borne by the remote processor.

In particular preferred embodiments of the present invention said local processor determines if an intervening update has occurred using a data key which accompanies said replacement version. This advantageously simple arrangement could take a number of forms, for example the data key may comprise a complete copy of the information indicating the files stored which the local processor originally passed to the remote processor, the local processor then being able to compare this with the current files stored. Alternatively, some sequence or letters or numbers could be passed between the processors. However, in preferred embodiments of the present invention said data key is indicative of how many times said local processor has been updated when said information was supplied to said remote processor and said local processor is adapted to compare said data key with how many times said local processor has been updated when said replacement version is received thereby determining if any intervening update has occurred.

Such an arrangement has the advantage that the amount of information passed between the processors is lowered thereby reducing the complexity of its interpretation and the time necessary for its transmission. The local processor need only com-

pare the number of times its files have actually been updated with the number indicated by the data key in order to determine if there has been an intervening update.

In a particularly simple and therefore advantageous embodiment of the invention, said local processor stores a number which is incremented each time said local processor is updated, said number being supplied by said local processor to said remote processor together with said information, said number being supplied back by said remote processor to said local processor together with said replacement version for comparison with said number currently stored by said local processor thereby determining if any intervening update has occurred.

As well as preventing an update if an intervening update has taken place, in preferred embodiments of the present invention said information is updated if replacement is not allowed in order that said remote processor may determine if replacement is still necessary.

This feature ensures that the remote processor with the most recent version of a file is allowed to update the local processor even if a remote processor holding an intermediate version of the file updates the local processor first.

In a particular preferred embodiment of the present invention said local processor stores flags to indicative of what information was supplied to each remote processor and upon receipt of a replacement version said local processor uses said flag for said remote processor supplying said replacement version to determine if an update has occurred since that remote processor was sent said information.

Using this feature the system can avoid having to pass information necessary to validate update requests between processors. Instead the local processor itself keeps track of what files it was storing when it supplied that information to a particular remote processor so that it may then check to see if any intervening updates have occurred before allowing an update from a particular remote processor to proceed. This arrangement has the advantage that the code for controlling the updating can reside solely on the local processor and thereby avoid the need to provide memory and processing for this function at each remote processor.

In an advantageously simple embodiment said local processor stores flags indicative of what information was supplied to each remote processor and upon receipt of a replacement version said local processor uses said flag for said remote processor supplying said replacement version to determine if an update has occurred since that remote processor was sent said information. When a request to update is received from a given remote

processor the local processor checks its current numerical indicator against the one recorded for that remote processor and only allows the update to proceed if they are the same. If the update is not allowed to proceed the local processor may as before supply the remote processor with the information concerning its current files and therefore allow the remote processor to determine if an update is still necessary.

Viewed from another aspect the present invention provides a method of operating a data processing system wherein at least one local processor stores a version of at least one file and a remote processor determines from information indicative of said version stored by said local processor if said version stored by said local processor should be replaced by a version of said file stored by said remote processor, characterised in that upon receipt of said replacement version said local processor determines if any intervening change has occurred in said version stored by said local processor since said information was supplied to said remote processor and allows replacement of said file if no intervening change has occurred.

An embodiment of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:

Figure 1 is a block diagram showing a data processing network of known type and to which the invention may be applied.

Figure 2 is a block diagram showing a host processor and a remotely connected personal computer in which the host can update data files within the personal computer in accordance with the present invention.

Figure 3 schematically illustrates the operation of an embodiment of the invention.

Figure 1 illustrates a data processing system in accordance with the present invention comprising a host processor 1 to which are connected a number of workstations. These workstations may be connected to the host processor 1 in a star configuration as represented by workstations 2 or through a controller controlled network as represented by controller 3 and workstations 4. Alternatively the workstation could be connected to the host processor, either locally as represented by loop 5 and workstations 6 or remotely as represented by loop 7 and workstations 8. In the latter case, the loop 7 may be a local area network (LAN) served by a server 9 connected to the host processor 1. One or more further host processors 10 and 10' may be provided to which some or all other various workstations 2, 4, 6 or 8 may be connected, either directly or indirectly.

Such data processing networks are known and no further description of the network per se is needed. The invention will be described in terms of

workstations, which are constituted by suitably programmed/configured personal computers such as the IBM Personal Computers, the IBM Personal Computer XT or the IBM Personal Computer AT or the IBM Personal System/2 personal computers. Generally, a personal computer will normally have a magnetic diskette file by means of which computer programs may be loaded on to it. It is possible to have a personal computer with no magnetic disc drives in which case any programming needs to be loaded over some kind of communications link.

In accordance with the present invention, use is made of a library of files stored within the host processor or processors and which is used to update files within a particular connected personal computer (workstation) when that particular workstation requires a particular file. There is no need for the host processor to keep a record of what files are kept by all the workstations which can be connected to it.

Figure 2 is a simplified block diagram showing the connection of a single host processor 11 to a workstation 12 over a communication link 13. As is normal, the workstation 12 constituted by a personal computer, includes a keyboard 14 and a display 15 operating under control of control logic 16 together with a store 17 for storing data and programs. Printers, plotters and other input/output devices, not shown, may also be connected to or form part of the workstation 12. The control logic 16 may be constituted by hardwired logic, programmed logic or combinations thereof. The store 17, normally a magnetic disk or diskette store, will contain one or more computer programs. The or each computer program will consist of a number of files. The files are variable in size represent an identifiable block of code consisting of one or more code modules. If an error is found within the file, that code either has to be replaced or modified to correct the error.

The host processor 11 therefore is provided with a store 18 which contains a library of object code which may be required by workstations connected to that host processor. If a file requires to be updated for any reason for a particular release/version of a PC program, then that file is updated. Once updated to the new level, the control logic 19 will ensure that any workstation 12 connected to the host processor 11 will be supplied with the new version file.

Figure 3 is a chart summarising the updating procedure. For the purposes of illustration, it is assumed that a particular PC program has two versions, version 1 stored on host 1 having files a, b and c at levels 1, 2 and 1 respectively and represented by a1, b2 and c1 and version 2 having data files a, b, c and d at levels 1, 2, 2 and 1

respectively stored on host 2. The object libraries of these two versions are stored in respective host processors. The personal computer has stored within it version 1 of the program with data files a, b and c all at level 1. In other words the version of the program installed in the personal computer is at a down-level. The updating process starts with one of the host processor sending a query to the PC.

As a result of the query from the host the PC responds with a reply which includes a list of the data files it has installed together with an indication of the levels of those files. Together with this indication the PC also sends the host processor a numeric key. the numeric key is a number held by the PC which is incremented each time the files stored by the PC are changed. Thus each value of the numeric key will be associated with a particular combination of files stored by the PC.

As contact with the workstation is made by any other host in the network, the same numeric key is returned to that host, since the specific combination of files and file levels is the same. If one of the hosts in the network determines that it holds files at a newer level than those held on the PC, it indicates to the PC that it is ready to perform an update operation. This indication is accompanied by the numeric key that was sent by the PC to the host at the time the file levels were reported. This forms a request for update from the host.

On receipt of the request for update the PC compares the numeric key of the request with the current numeric key of the PC. If the values of the keys match, then the host is reacting to the same combination of files and levels as currently exist on the PC and the pending update operation is valid. The PC increments the numeric key, and grants the host permissions to continue. The PC records that an update is now in progress and sets an update in progress flag. However, if the values of the keys differ then the files or levels on the PC have been updated, possibly by another host or by the PC, in the interval since the files and levels were reported. In this case the pending update operation by this host may be invalid since the library file level may no longer be more recent than that of the PC. The PC denies the requesting host permission to update, and returns instead the current file list and levels, together with the current numeric key. The host may then re-evaluate the need for update, and, only if still required, repeat the update request.

While a valid update is in progress, a further host may make a request for update. At this time, the PC should either immediately deny permission, or queue the request to be dealt with when the current update process is complete, at which time the request would be denied and the new current

files, levels and numeric key would be returned.

When any valid update completes, the PC resets the update in progress flag, to allow further update by other hosts, subject to the numeric key validation process described.

An alternative embodiment of this invention may not actually send the numeric key from the PC to the host, but may instead record the key in the PC together with the host to which it has been sent. When an update request is received from a particular host, the PC validates the key held in the PC for this host against the current value of the key, using the rules described above. This avoids the need for the host to receive and resend the key.

Returning now to the example shown in Figure 3, both host 1 and host 2 query the PC concerning the files it is storing and their levels. The PC then returns this information to both hosts together with the numeric key having value n . Both hosts then analyse the information sent them to determine if an update is required. Both hosts determine that an update is required. For some reason, such as faster processing or communication, host 1 makes its request for the update before host 2. The request for update from host 1 causes the update in progress flag to be set which suspends the update being attempted by host 2.

Together with the request for update host 1 returns the numeric key n to the PC. The PC compares this to the current value of the numeric key and determines that the update is valid. The PC then increments the numeric key to a value $n + 1$ and allows the update to proceed thereby replacing files $a1, b1, c1$ with files $a1, b2, c1$. After the update from host 1 has finished the PC resets the update in progress flag which allows the update being attempted by host 2 to be processed further.

The request for update from host 2 is accompanied by the numeric key n and this time when the PC compares the numeric key being sent it with the current value of $n + 1$ it determines an intervening update has taken place.

The PC refuses the request for update and sends host 2 its new list of files and versions together with the new numeric key so that host 2 may determine if an update is still required. In this case host 2 holds more recent files than host 1 so an update will still be required and host 2 will carry out this update in the manner described above for host 1. Host 2 will thereby replace files $a1, b2, c1$ with files $a1, b2, c2, d1$.

Each host does not have to keep a record of what programs and their levels that are installed in the workstations connected to it. Upgrading of the version stored in the host will ensure upgrading of the versions stored in the connected PCs as they are needed.

The invention is applicable to PC programs which communicate with hosts. However, it is possible to provide this facility for any PC program installed on a PC provided that the PC can communicate with a host computer which has a copy of that program's data files within the object library.

Where the host is able to recognise what version of a program is installed at the personal computer, the mechanism described can also be used to supply a data file to the PC should the PC's response indicate that it does not contain a data file which should be included in that version of the program.

Claims

1. A data processing system having at least one local processor storing a version of at least one file and a remote processor determining from information indicative of said version stored by said local processor if said version stored by said local processor should be replaced by a version of said file stored by said remote processor, characterised in that upon receipt of said replacement version said local processor determines if any intervening change has occurred in said version stored by said local processor since said information was supplied to said remote processor and allows replacement of said file if no intervening change has occurred.

2. A data processing system as claimed in claim 1, wherein said local processor supplies said information to said remote processor.

3. A data processing system as claimed in claims 1 or 2, wherein said data processing system includes a plurality of said remote processors.

4. A data processing system as claimed in claims 1, 2 or 3, wherein said local processor determines if an intervening update has occurred using a data key which accompanies said replacement version.

5. A data processing system as claimed in claim 4, wherein said data key is indicative of how many times said local processor had been updated when said information was supplied to said remote processor and said local processor is adapted to compare said data key with how many times said local processor has been updated when said replacement version is received thereby determining if any intervening update has occurred.

6. A data processing system as claimed in claim 5, wherein said local processor stores a number which is incremented each time said local processor is updated, said number being supplied by said local processor to said remote processor together with said information, said number being supplied back by said remote processor to said local processor together with said replacement ver-

sion for comparison with said number currently stored by said local processor thereby determining if any intervening update has occurred.

7. A data processing system as claimed in claims 1, 2 or 3, wherein said local processor stores flags indicative of what information was supplied to each remote processor and upon receipt of a replacement version said local processor uses said flag for said remote processor supplying said replacement version to determine if an update has occurred since that remote processor was sent said information. 5 10

8. A data processing system as claimed in claim 7, wherein said local processor stores a number which is incremented each time said local processor is updated and said flags are indicative of the value of said number when said information is supplied. 15

9. A data processing system as claimed in any preceding claim, wherein said information is updated if replacement is not allowed in order that said remote processor may determine if replacement is still necessary. 20

10. A method of operating a data processing system wherein at least one local processor stores a version of at least one file and a remote processor determines from information indicative of said version stored by said local processor if said version stored by said local processor should be replaced by a version of said file stored by said remote processor, characterised in that upon receipt of said replacement version said local processor determines if any intervening change has occurred in said version stored by said local processor since said information was supplied to said local processor and allows replacement of said file if no intervening change has occurred. 25 30 35

40

45

50

55

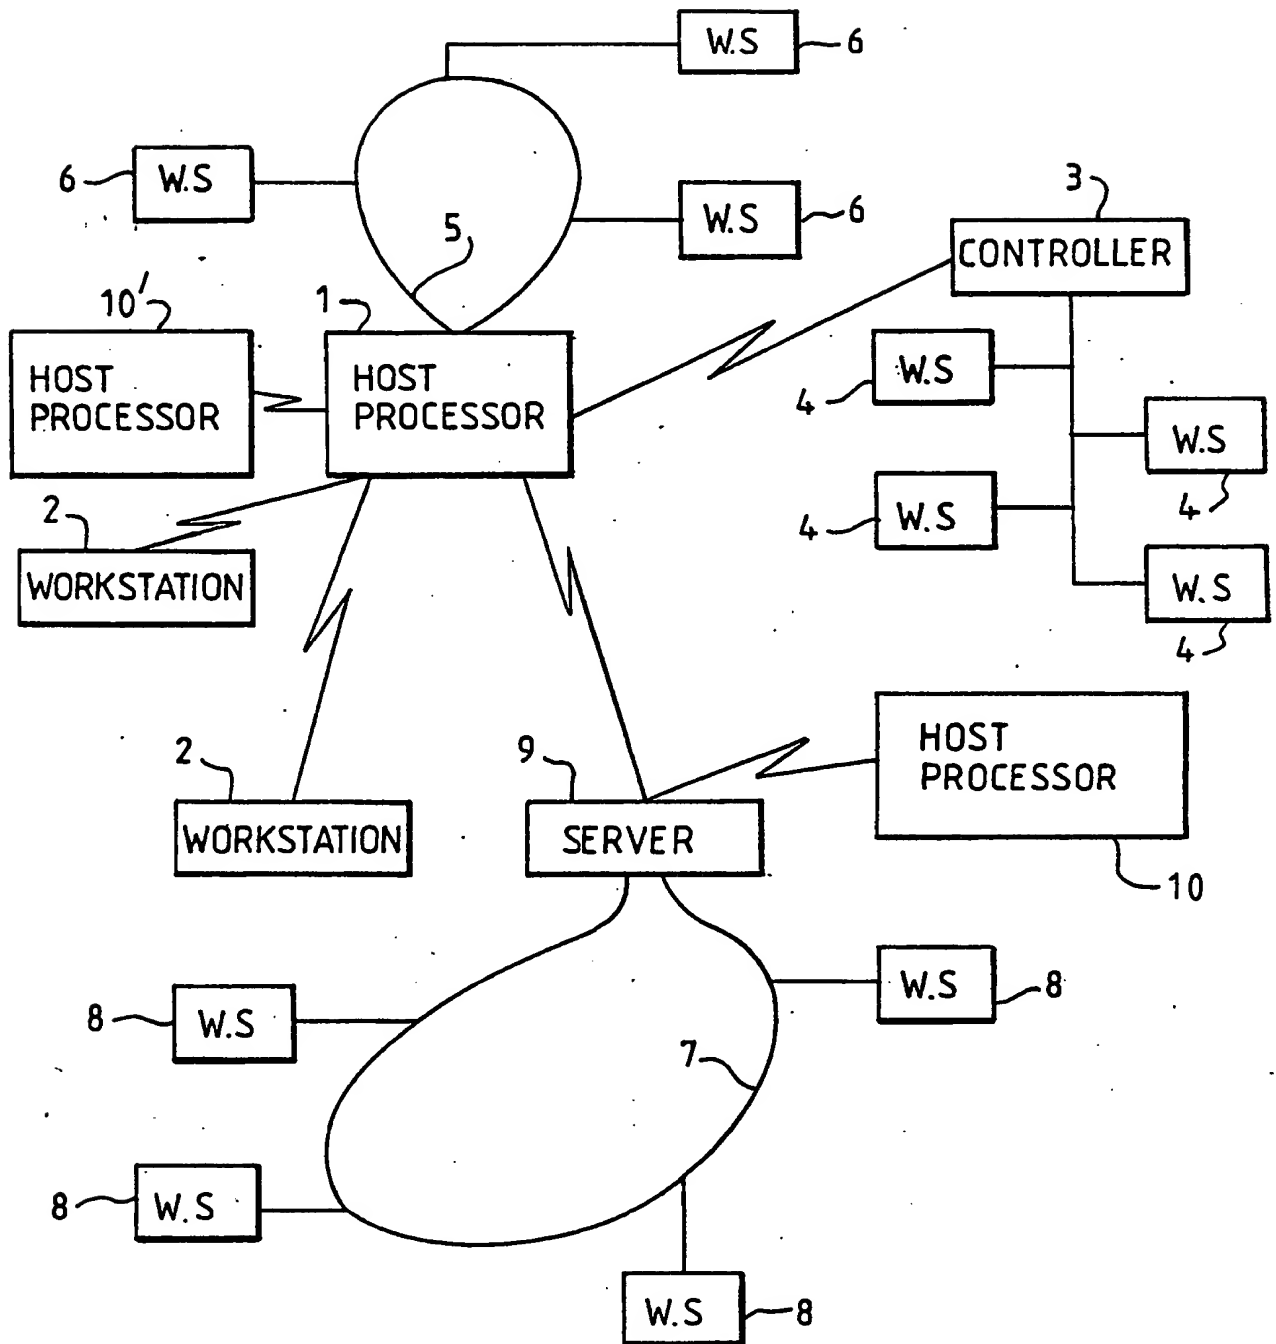


FIG. 1

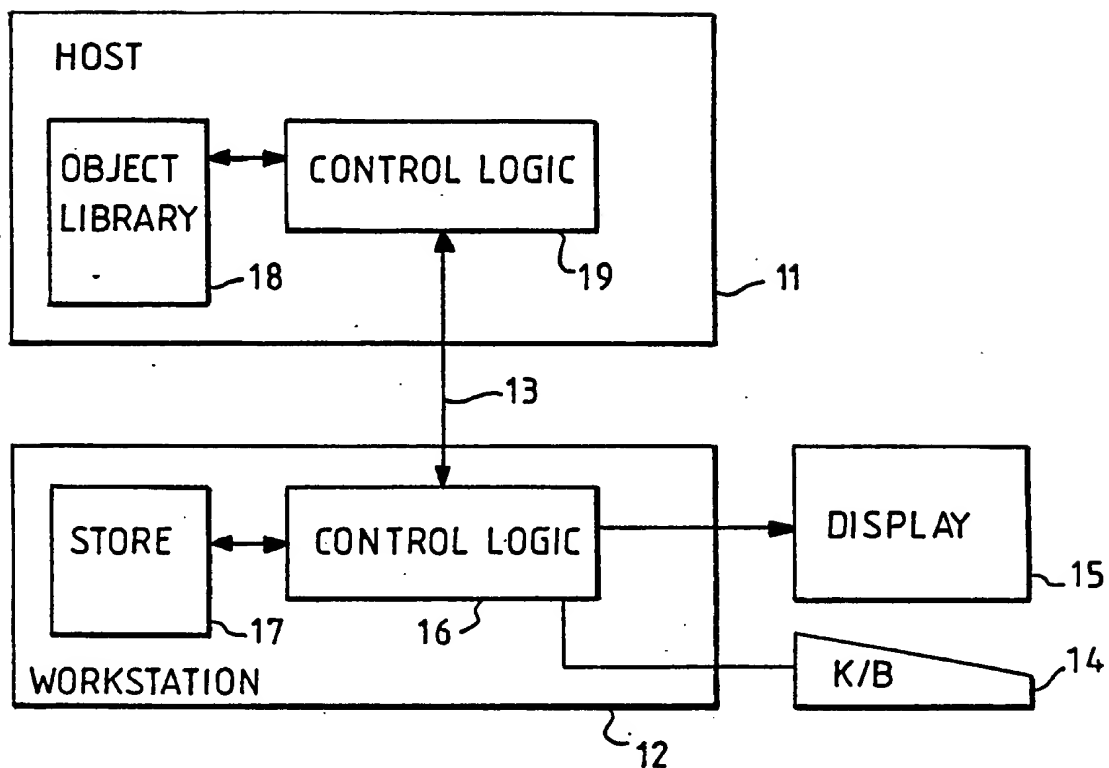


FIG. 2

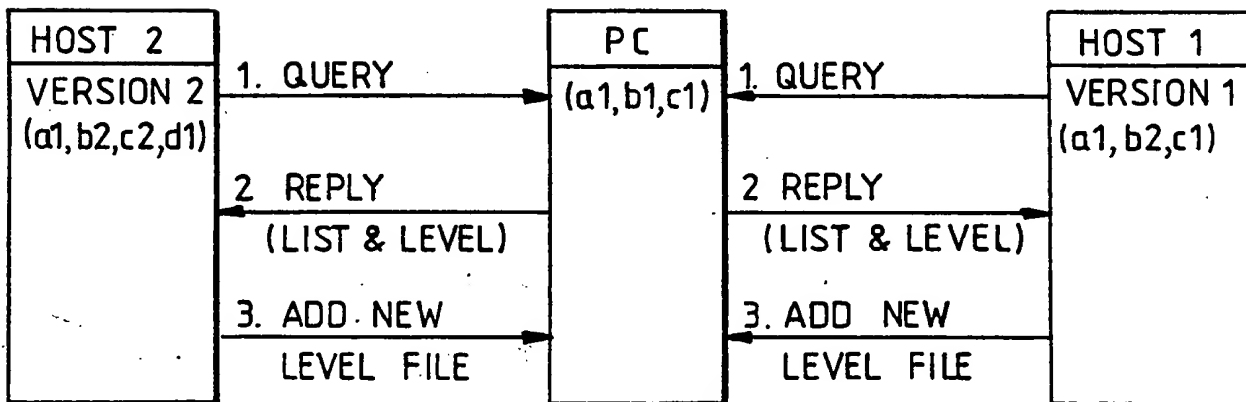


FIG. 3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 89 30 4816

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl.5)
D,Y	EP-A-0 284 924 (IBM) * Column 2, lines 26-49; column 4, lines 2-13,37-56; figures 1,3 * ---	1-6,10	G 06 F 15/403 G 06 F 9/44
Y	IEEE TRANSACTIONS ON SOFTWARE ENGINEERING, vol. SE-13, no. 1, January 1987, pages 39-47, IEEE, New York, US; S.K. SARIN et al.: "Discarding obsolete information in a replicated database system" * Page 39, left-hand column, lines 12-20; page 39, right-hand column, line 37 - page 40, left-hand column, line 21; page 40, left-hand column, line 48 - right-hand column, line 3 * ---	1-6,10	
A	EP-A-0 290 828 (HITACHI) * Abstract; column 6, line 35 - column 7, line 46; column 8, lines 16-32 * ---	1-6,10	
A	IEEE COMPUTER SOCIETY, OFFICE AUTOMATION SYMPOSIUM, Gaithersburg, MD, 27th-29th April 1987, pages 171-177, IEEE, New York, US; S.H. SON: "A synchronization scheme for replicated data in distributed information systems" * Sections 2.2,5.2 * -----	1,10	TECHNICAL FIELDS SEARCHED (Int. Cl.5) G 06 F
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 21-12-1989	Examiner WILTINK J.G.
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document I : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 96/01494

A. CLASSIFICATION OF SUBJECT MATTER

IPC 6 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category of documents with indication, where appropriate, of the relevant passages

Relevant to claim No.

P,X	US,A,5 434 994 (SHAHEEN AMAL A ET AL) 18 July 1995 see abstract see column 1, line 1 - column 3, line 17 ---	1,20,39, 50
A	EP,A,0 278 315 (IBM) 17 August 1988 see page 2, line 1 - page 4, line 47 ---	1,20,39, 50
A	US,A,4 853 843 (ECKLUND DENISE J) 1 August 1989 see abstract see column 1, line 1 - column 3, line 64 ---	1,20,39, 50
A	EP,A,0 396 834 (IBM) 14 November 1990 see abstract; figure 1 ---	1,20,39, 50

-/--

☒ Further documents are listed in the continuation of box C.☒ Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- * & * document member of the same patent family

Date of the actual completion of the international search

4 June 1996

Date of mailing of the international search report

14. 06. 96

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epó nl,
Fax (+31-70) 340-3016

Authorized officer

Katerbau, R